

REMARKS

This application was examined with claims 51 through 91. No claims are allowed. No claims are canceled. Claims 51, 53 through 55, 60, 61 and 76 are amended. Claims 51 through 91 remain in the application.

Applicant requests reconsideration and reexamination of the above-identified application in view of the amendments made to the specification and claims. The following remarks state Applicant's bases for making this request and are organized according to the Examiner's Action by paragraph number.

Examiner's Action, Paragraphs 2 through 4

The Examiner rejects the claims under 35 U.S.C. 101 as being related to non-statutory matter because the claims are not limited to tangible embodiments. Applicants are amending each of independent claims 51, 60, and 76 to indicate the program is "for storage in a computer readable medium". Applicants respectfully submit that these amendments overcome the rejection as each claim now is directed to a tangible embodiment.

Examiner's Action, Paragraph 5

The Examiner objects to incorrect dependencies in claims 53 and 61. Claim 53 is amended to depend from claim 51; claim 61 from claim 60.

Examiner's Action, Paragraphs 6 and 7

The Examiner rejects the application on the basis of non-statutory double patenting rejection based on the judicially-created doctrine. Applicants concurrently filed Terminal Disclaimer is believed to overcome this rejection.

Claim Rejections - Prior Art

The Examiner rejects all the claims on the basis of two references; namely: U. S. Patent No. 6,324,654 to Wahl et al. for a Computer Network Remote Data Mirroring System (the "Wahl reference") and U. S. Patent No. 6,085,298 to Ohran for Comparing Mass Storage Devices Through Digests That Are Representative of Stored Data in Order to Minimize Data Transfer (the "Ohran reference"). Before discussing each of the Examiner's rejections in detail, it may be helpful to review what Applicants regard as their invention.

More specifically, referring to FIG. 1 of the application, the claimed invention is particularly adapted for use systems with remote backup facilities, such as the second site 23, in which the communications link 36 can be tailored for the particular backup requirements. Parameters of importance include:

- (1) the amount of data to be transferred over a communications link;
- (2) the number of individual data paths in the communications link;
- (3) the bandwidth of each data path; and
- (4) the time between the beginning of an update to the remote or target backup and its resynchronization with the source.

Using certain data, an administrator can accurately predict various bandwidth requirements for a link, such as the communications link 36.

In accordance with Applicants' invention a program, such as the collector and reporter in memory locations 50 and 52 in the host 24 in the specifically disclosed embodiment, the claimed program determines from activities at the local production site the requirements for the remote data communications link 36. The administrator can then manipulate the various parameters for achieving a maximum advantage. For example, if the amount of data to be transferred and resynchronization times are known from actual measurements, the reporter can provide a prediction of the required bandwidth for communications link. This in turn can determine the bandwidth for parallel data paths in the communications link, should they be available. Alternatively, for a given amount of data and

bandwidth, the administrator can predict the resynchronization time that will result.

Claim 51 defines a step of defining a data group set and an interval. This is accomplished by means of a collection command as described at col. 6, lines 4-36. When that command is received, the program monitors a first update to each data block location in the data group set. For example, in the application the program allocates a PB bit table for a cylinder block such that one PB bit corresponds to each specific track defined in a data group set as established in a configuration table in FIG. 2. Each PB bit records a first write or update to a corresponding data track. Subsequent updates to that data track do not alter the state of the corresponding PB bit.

Once the interval for acquiring data has been completed, the information transfers to an entry in the data set 51 in FIG. 1. From this the reporter 52 determines the total number of data block locations that were updated at least once during the defined interval. For example, for a controller summary report, as shown in FIG. 4, steps 97, 100 and 102 logically OR all the entries for the time interval. This determines the total number of data block locations that were altered during the interval. This data then can be converted into bandwidth-based information about all the transfers during the defined interval. The bandwidth-based information is dependent upon the particular information being sought but includes the basic

data communications link parameters from which the previously described predictions can be made.

In the Office Action the Examiner references specific sections of the Wahl and Ohran references. Exhibit A consolidates these referenced sections for use in the following remarks. Exhibit A presents four columns. The "Referenced Section" column uses "W" to designate the Wahl reference; "O", the Ohran reference. The following number is an identifier. The "Claim" column identifies each claim of claims 51-75 against which the Examiner cited the corresponding Referenced Section. "Col/Line Ref" column contains a citation; the last column, the referenced text. In some situations the cited text incorporates language in addition to the language in the Examiner's cite to incorporate a complete sentence in the citation and referenced text; the citation is modified accordingly. In the following argument, references are to the Referenced Section. For example, cite "Referenced Section W1" means to refer to the text at column 3, lines 22-28 in the Wahl reference. Exhibit A also indicates that this text of Referenced Section W1 was only applied to claim 51. It is hoped that Exhibit A will be of benefit to Applicants' arguments for allowance.

Examiner's Action, Paragraphs 8 and 9

The Examiner rejects claims 51 through 59 under 35 U.S.C. 102(e) as being anticipated by the Wahl reference.

Applicants respectfully traverse the rejection of claims 51 through 59 on the basis of the Wahl reference.

The Wahl reference discloses a computer network with remote data mirroring. A "write" to a local data device also produces a write to a local chronologically sequenced journal stage entry or "writelog". Memory space is assigned to the writelog device to prevent memory overflow. Remote mirroring occurs when logical groups of local data device and writelog device pairs are supplemented with a primary mirror daemon that communicates to a secondary system. Each change is then committed to the remote computer system. The Wahl reference enables an administrator to produce "throttles" that enable a predetermined portion of network bandwidth to be assigned to the mirroring based on user-selected criteria.

While the Wahl reference discloses an approach for providing data mirroring at a remote location, it does not disclose a method, apparatus or program for allowing an administrator to predict future operations whereby the administrator can select the values for the different communications link parameters, such as data path bandwidth, the number of parallel data paths in a link and resynchronization times.

Claim 51 defines a program for providing information about the number of updates to data blocks that is useful in predicting bandwidth-based information about those updates. Certain elements of the claim, read independently of the remainder of the claim, may define processes that are in the prior art. However, the Wahl reference does not disclose or suggest the claimed combination of those elements. For example, the Examiner argues that Referenced Section W1 and W2 disclose Paragraphs A) and B), respectively. Referenced Section W1 only defines the basic concepts of grouping volumes into logical groups. However, this grouping is not in the context of the claimed program. Similarly, with respect to Paragraph B).

With respect to Paragraph C) the Examiner relies on Referenced Section W3 to anticipate a recording process in which only a first update to each data block location is recorded. However, the writelog device in the Wahl reference records each and every change to the data. Metadata in the writelog device is updated every  $m^{th}$  time an entry is written. This is not the operation of Paragraph C).

With respect to Paragraph D), Referenced Section W3 updates metadata. This is not a transfer of information obtained from the recording that is a list of each of the data blocks for which at least one update has been made to a data group set of information.

The Examiner also relies on Referenced Section W3 as disclosing the step of determining the total number of data block locations that were updated at least one time during the defined interval. Applicants do not see how that language could disclose or suggest the substance of Paragraph E) to a person of ordinary skill in the art. The Examiner uses Referenced Sections W4 and W5 to support his rejection of the conversion of Paragraph F). Applicant respectfully submits that nothing in these referenced sections, that describe basic computer concepts, or elsewhere in the Wahl reference discloses or suggests the conversion of Paragraph F).

Consequently, Applicants respectfully submit that claim 51 should be allowable over the Wahl reference.

The Examiner rejects claims 52 and 53 on the basis of Referenced Section W6. Each of claims 52 and 53 relate to a more specific definition of the recording of a first update to a data block location. Referenced Section W6 describes operations that occur with respect to a last entry and the setting of a flag to request acknowledgement. It is respectfully submitted that this operation in Referenced Section W6 does not disclose or suggest the concept set forth in either claim 52 or 53. Therefore claims 52 and 53 should be allowable.

Claim 54 is rejected on the basis of Referenced Sections W7, W8 and W9. Referenced Section W7 describes data mirroring



for providing volume grouping or logical groups. Referenced Section W8 refers to the use of time stamps on each entry to the writelog device 18. Referenced Section W9 states that the writelog device 18 can comprise a disk drive device.

Applicants see nothing in these referenced sections that describes the storage of a time interval, data group and defined flags in another data set as set forth in claim 54.

The Examiner rejects claim 55 on the basis of Referenced Sections W10 and W6 and rejects claim 56 on the basis of References Section W6. Referenced Section W10 discloses the operation of a computer network system in a mirrored mode. Referenced Section W6, as previously stated, discusses setting a flag when the last data entry of an internal buffer is prepared to be sent. Applicants see nothing in these disclosures that anticipates or suggests Applicants' invention as defined in either claim 55 or claim 56.

The Examiner rejects each of claims 57, 58 and 59 on the basis of Referenced Sections W11 and W12. Referenced Section W11 describes how the Wahl reference method and apparatus meet the various objectives of that invention. Referenced Section W12 refers to an operational mode in which a coherent copy of data exists in a secondary computer system. Applicants respectfully submit that neither referenced section has anything to do with the generation of a report such as described generally or specifically in either claim 58 or 59.

Examiner's Action, Paragraphs 10 and 11

The Examiner rejects claims 60 and 76 under 35 U.S.C. 102(e) as being anticipated by the Ohran reference. Specifically, the Examiner argues that Referenced Section 01 discloses all the information in the preamble, Paragraph A) and Paragraph E) of claims 60 and 76; Referenced Section 02 Paragraph B); Referenced Sections 03 and 04, Paragraph C); and Referenced Section 05, Paragraph D).

Applicants respectfully traverse this rejection.

As Applicants understand the Ohran reference, a backup system transfers only changed records to a backup site. The system assumes that two storage systems, such as disk storage systems at the primary and backup sites, are initially synchronized. When a record in the primary system changes, the storage location is identified. A snapshot of the primary is taken when the primary is logically consistent. This way a backup is created of only the changed data to reduce the impact on bandwidth between the primary and backup sites.

Although the Ohran reference discloses a backup system with some similarities to the backup system disclosed in Applicants' application and U. S. Patent No. 6,662,197, the Ohran reference does not disclose any monitoring of operations for the purpose of predicting bandwidth related properties of the communications path and related parameters. Therefore

Applicants respectfully submit that independent claims 60 and 76 are directed to an invention that is entirely different from the disclosure of the Ohran reference.

Further, referenced Section 01 generally discusses the objectives of the Ohran reference in terms of transporting a last updated version of a record to the backup site at a time when the primary device is in a logically consistent state. Again, while recognizing this can reduce the required bandwidth requirement, the Ohran reference does not suggest prediction of those requirements based on actual usage. Referenced Section 02 discusses the initiation of backups based upon a combination of time interval and logically consistent backup state from one or more primary systems. As previously indicated, the elements of Paragraph A) and B) are broadly recited and could read on these particular elements as disclosed in the Ohran reference. However, the combination of all elements present a patentable claim as will now be discussed.

In accordance with claims 60 and 76, only the identification of each defined data block that is updated is recorded. Referenced Sections 03 and 04 describe process of obtaining a snapshot which records the data blocks themselves.

Applicants respectfully submit this does not meet the requirements of Paragraph C).

With respect to Paragraph D) the Examiner relies on Referenced Section 05. In Paragraph D) a determination of the

number of updated defined data blocks during the time interval is made. Referenced Section 05 merely provides alternative definitions of "data block". Applicants see nothing in this section that discloses or suggests the "determination" of Paragraph D).

Consequently, Applicants respectfully submit that claim 60, that is directed to a program for allowing the determination of bandwidth related characteristics for a communications path that extends between first and second remote sites even from a primary site from which the updates are generated, is not disclosed or suggested in the Ohran reference.

As claim 60 defines the data sets in terms of defined data blocks and as Claim 76 defines the data sets in terms of data blocks comprising disk tracks, Applicants respectfully request the Examiner to withdraw his rejection of claims 60 and 76.

Examiner's Action, Paragraphs 12 and 13

The Examiner rejects claims 61 through 75 and 77 through 91 under 35 U.S.C. 103 as being unpatentable over the Ohran reference in view of the Wahl reference. Each rejection with respect to claims 61 through 75 relies on disclosures in particular referenced sections. Claims 77 through 91 are rejected for the same reasons claims 61 through 75 are rejected, respectively.

Applicants respectfully traverse of each of these rejections.

With respect to claims 61 and 77, the Examiner relies upon Referenced Sections 03, 04 and W13. As previously argued, Applicants do not believe that Referenced Sections 03 and 04 disclose or suggest the table of claim 61, subparagraph i) or ii). Applicants do not believe that Referenced Section W13, describing the organization of a data storage unit, relates to the "table" of claim 61, subparagraph i). Further, Applicants respectfully submit that nothing in the references, including the referenced sections would lead a person of ordinary skill in the art to reach the Examiner's stated conclusions concerning remote data mirroring systems configured for optimal data mirroring and the problem of limited bandwidth. The Examiner comments that Referenced Section 06 discloses recording over a plurality of time intervals and recording sets of all the table positions to the first state at the beginning of each time interval. Referenced Section 06 refers to write operations merely in terms of writing to a mass storage device and the inclusion of means for performing that operation. The backup discussion of Referenced Section 02 defines a concept of conventionally implemented as differential and/or periodic backups. It is not related to the same issue as the time intervals of Applicants' claimed invention.

Claims 63 and 79 introduce the concept of multiple recordings for successive time intervals thereby to provide additional information from which the bandwidth based information can be obtained. Referenced Section O7 is directed to a data shadowing program in which data is written to a file and backed up each time. This is not the concept of recording the number of accesses during each of successive intervals with a date-time stamp.

With respect to claims 64 and 80, Referenced Section W14 discloses some commands that relate to several operations including "throttles". The broad concept of defining a report to include particular parameters is old. However, in combination, the elements of claims 64 and 80 define a specific report that would not be obvious in view of Referenced Sections O7 and W14.

Claims 65 and 81 specifically relate to a controller report wherein all the updates for all logical volumes controlled by the controller are summed. Referenced Section W15, upon which the Examiner relies, describes the organization of logical volumes into logical groups. Nothing in this referenced section, however, suggests the controller report of claims 65 and 81.

Claims 66 and 82 further define the report of claims 65 and 81, respectively, by defining a logical OR operation. Referenced Sections W13 and W16, upon which the Examiner

relies, do not disclose or suggest the concept of these claims. Referenced Section W13 has been discussed previously. Referenced Section W16 merely refers to a choice of a relational or transition logical operator.

Claims 67 and 83 further define the controller report of claims 65 and 66 and claims 81 and 82 in which a resynchronization time parameter is established as a given requirement so that the processing provides the total required bandwidth of the communications path to provide the required resynchronization time. None of Referenced Sections W12, W14, W17 or W18 discloses or suggests this resynchronization. Referenced Sections W12 and W14 have previously been discussed.

Referenced Section W17 generally discloses the concept of selecting bandwidth, CPU or writelog device utilization for control. Referenced Section W18 describes, asynchronous, synchronous or near synchronous data mirroring. This is the different from the resynchronization time of claims 67 and 83. Therefore Applicants respectfully submit that none of the referenced sections suggests the combinations of claims 67 and 83.

Claims 68 and 84 further define the report and a particular method for determining bandwidth where multiple paths make up a communication link. The Examiner relies on Referenced Sections W4 and W18. There is a description of bandwidth throttling in Referenced Section W4 and the concept

of ongoing mirroring and maintenance of a coherent copy in a secondary computer system in Referenced Section W18. Neither of these referenced sections, however, disclose or suggest the subject matter of claims 68 and 84.

Similarly, claims 69 and 85 define a specific implementation of the program by determining the number of required parallel data paths of a given bandwidth to achieve a required bandwidth. Applicants respectfully submit that these claims are patentable for the same reason as claims 68 and 84 are patentable over Referenced Sections W4 and W18.

Claims 70 and 71 and claims 86 and 87 refer to alternate implementations of the reporter command response whereby the program uses the available bandwidth of communications path to determine a resynchronization time. Applicants respectfully submits that the Referenced Sections W4, W14 and W18 do not render claims 70, 71, 86 and 87 obvious for reasons that have been described previously.

Claims 72 and 73 define still another variation on the reporter operation. Particularly, these claims reference the concept of maintaining the updated data block information on a logical volume basis. Applicants see nothing in Referenced Sections W13, W14 or W16 that disclose or suggest any such report.

Claims 74 and 75 refer to program features that produce another possible report. In Claims 74 and 75 the report is



based on data sets whereby all the data block changes are combined for a single data set in claim 74 and a final table that is the logical OR of all the data set tables in claim 75. Applicants see nothing in Referenced Sections W13, W14 or O8 that disclose or suggest the concepts of claims 74, 75, 90 or 91.

#### Summary

Applicants have amended claims 51, 53 through 55, 60, 61 and 76. Applicants respectfully submit that each of claims 51 through 91 now are directed to statutory subject matter. Applicants concurrently filed Terminal Disclaimer should overcome the double-patenting rejection. Applicants have further presented specific arguments as to why each of claims 51 through 91 should be allowed over the prior art, particularly the Examiner's Referenced Sections as shown in Exhibit A. Consequently, Applicants respectfully request the Examiner to withdraw the rejections and to allow the application as amended with claims 51 through 91.

If there are any questions, we urge the Examiner to call us collect.

Respectfully Submitted,

/George A. Herbster/

GEORGE A. HERBSTER, Attorney  
40 Beach Street  
Manchester, MA 01944  
(978) 526-8111

## EXAMINER'S REFERENCED SECTIONS

US PATENT NO. 6,324,654 – WAHL REFERENCE			
REFERENCED SECTION	CLAIM	COL/LINE REF	WAHL REFERENCE TEXT
W1	51	3/22-28	Also in accordance with the present invention, the computer network remote data mirroring system can be structured to provide volume grouping, or logical groups. Consequently, data at the local, or primary, site can be replicated at a plurality of remote sites, as compared to known architectures which provide point-to-point (local to a single remote site) data mirroring.
W2	51	7/25-29	A reserved area at the beginning of the writelog device <b>18</b> , that contains metadata about the write-log device, is updated every $m^h$ time an entry is written to or read from the writelog device or if a configurable period of time has elapsed since the last update of the metadata.
W3	51	7/18-29	Each entry written to the writelog device <b>18</b> consists of data and a header. The header contains important information used by other system components, such as a timestamp, sequence number, device offset, and size of the transaction. The oldest data entries are read from the end or tail of the writelog device <b>18</b> and sent across the network <b>20</b> , while new data entries are written to the beginning or head of the writelog device. A reserved area at the beginning of the writelog device <b>18</b> , that contains metadata about the write-log device, is updated every $m^h$ time an entry is written to or read from the writelog device or if a configurable period of time has elapsed since the last update of the metadata.
W4	51, 68, 69, 71	3/39-44	The computer network remote data mirroring system of the present invention additionally provides network band-width throttling. Bandwidth throttling enables a predetermined portion of the network bandwidth to be assigned to remote data mirroring depending on the time of day or other criteria.
W5	51	5/35-38	Both the primary and secondary computer systems are configured to have adequate amounts of disk storage and network bandwidth allocated to accommodate the flow of data needed to provide remote data mirroring.
W6	52,53, 55,56	21/20-25	When the last data entry of the internal buffer for a local data storage unit <b>26</b> is being prepared to be sent, a flag is set in the header requesting an acknowledgment from the child remote mirror daemon <b>30A</b> , <b>30B</b> once the data of the entry has been committed to the mirror device <b>32</b> .
W7	54	3/22-24	Also in accordance with the present invention, the computer network remote data mirroring system can be structured to provide volume grouping, or logical groups
W8		7/18-29	Each entry written to the writelog device <b>18</b> consists of data and a header. The header contains important information used by other system components, such as a timestamp, sequence number, device offset, and size of the transaction. The oldest data entries are read from the end or tail of the writelog device <b>18</b> and sent across the network <b>20</b> , while new data entries are written to the beginning or head of the writelog device. A reserved area at the beginning of the writelog device <b>18</b> , that contains metadata about the write-log device, is updated every $m^h$ time an entry is written to or read from the writelog device or if a configurable period of time has elapsed since the last update of the metadata.
W9		7/ 50-53	In the embodiment in which the writelog device <b>18</b> comprises a disk drive device, additional disk storage is dynamically assigned or another disk storage device is chained into the primary computer system <b>12</b> .

US PATENT NO. 6,324,654 – WAHL REFERENCE			
REFERENCED SECTION	CLAIM	COL/LINE REF	WAHL REFERENCE TEXT
W10	55	5/28-38	The computer network remote data mirroring system of the present invention accomplishes this result through time-sequenced transfers of data from the primary computer system to the secondary computer system over the network. Should a failure occur on the primary computer system, the secondary computer system can provide immediate access to contemporary application data. Both the primary and secondary computer systems are configured to have adequate amounts of disk storage and network bandwidth allocated to accommodate the flow of data needed to provide remote data mirroring.
W11	57,58, 59,67	4/35-45	The computer network remote data mirroring system of the present invention achieves high application performance by implementing asynchronous, synchronous, or near synchronous data mirroring using network bandwidth throttling. It provides substantially real-time data mirroring over LANs and WANs to quickly move data offsite, yet does not impact application performance significantly. In the event of a disaster taking the primary data center out of service for hours or days, operations can be transferred to a secondary site within minutes, operating on an up-to-the-minute copy of the original data set.
W12	57,58, 59, 67	18/40-44	This is the operational mode in which remote data mirroring is accomplished and in which ongoing mirroring and a coherent copy of data exists on the secondary computer system 14.
W13	61,66, 73	8/21-26	As shown in FIG. 3, each local data storage unit <b>26</b> appears as a raw disk partition to the kernel. Therefore, each local data storage unit <b>26</b> accepts and handles any request that can be made to a normal raw disk partition or fixed size volume, such as create and mount a file system, or support DBMS table space allocations.
W14	64, 67, 70, 72, 74	22/34-49, 54-60	<p>The qds info command generates an ASCII report for one or more local data storage units <b>26</b> from the perspective of the device driver 22 and the writelog device <b>18</b>. The qds info command indicates if the writelog device <b>18</b> has been placed in a special mode, for example, the bypass mode or the refresh mode. It also shows performance metrics specific to the writelog device <b>18</b>.</p> <p>qdsconfigtool is a graphical user interface utility for viewing, editing, or defining logical group configuration files, including primary and secondary computer systems 12, 14, tunable primary mirror daemon parameters, writelog device extension pools <b>18A</b>, local data storage units <b>26</b>, and throttles.</p> <p>For example, FIG. 12 illustrates a chart generated by qdsperftool. Additionally, the user can display multiple charts at one time, modify the charts, delete them or print them. qdsperftool enables the user to observe performance of the computer network remote data mirroring system <b>10</b> over time and shows trends.</p>
W15	65	11/44-54	Also in accordance with the present invention, the computer network remote data mirroring system <b>10</b> can be structured to provide volume grouping, or logical groups 34, as shown in FIG. 5. That is, a collection of local data storage units <b>26</b> can be configured as a coherent unit, called a logical group 34. In one exemplary implementation, the computer network remote data mirroring system <b>10</b> supports up to <b>512</b> logical groups 34, each with an unlimited number of local data storage units <b>26</b> . Placing affiliated local data storage units <b>26</b> in the same logical group 34 is an effective way to configure an efficient system.
W16	66, 73	15/58-61	Fourth, the user chooses a relational or transitional logical operator from the pulldown menu in the center of the screen, as shown in FIG. 8.

US PATENT NO. 6,324,654 – WAHL REFERENCE			
REFERENCED SECTION	CLAIM	COL/LINE REF	WAHL REFERENCE TEXT
W17	67	15/50-52	The three options determine which system component, namely, network bandwidth, CPU, or writelog device utilization, is being controlled.
W18	68,69,71	16/49-63	<p>The dynamic controls of the computer network remote data mirroring system <b>10</b> enable the user to define how much and when computer system/network resources are used for remote data mirroring.</p> <p>An example of network bandwidth throttling is illustrated in FIG. <b>10</b>. In the example, the following throttles regulate network bandwidth consumption. The first two throttles deal with maintaining usage of the network <b>20</b> below a certain point. Note that "sleep" is incremented by 15,000 micro-seconds if usage exceeds 200 KB per second. If network usage continues to increase and exceeds 300 KB per second, "sleep" is incremented by 5,000 microseconds every time the throttle evaluates true. The remaining throttles focus on maintaining network usage. If usage begins to decline, "sleep" is decremented continuously until it reaches zero.</p>

US PATENT NO. 6,085,298 – OHRAN REFERENCE			
REFERENCED SECTION	CLAIM	COL/LINE REFERENCE	OHRAN REFERENCE TEXT
O1	60, 60A, 60E, 76	6/25-39	Because the present invention takes a data block approach to the backing up of a mass storage system, the present invention minimizes the amount of data that needs to be transferred to make a backup to the absolutely minimum possible. For example, if a large database has five records that change, prior art systems would copy the entire large database. The present invention, however, copies only the five records that have changed. Because the amount of data is minimized, the present invention is particularly well suited to backing up data to a backup system located at a remote site. The present invention can utilize low bandwidth communication links to transfer backup data to a remote backup site. As an example, in many cases conventional dial-up telephone lines with a 56.6 k baud modem will be entirely adequate for many situations.
O2	60B, 62	21/57-62	The backups may be initiated, either by the backup system or by the primary systems, on a periodic schedule. Thus, the present invention may be used to capture a series of backups, each representing a logically consistent backup state, from one or more primary systems.
O3	60C, 61	17/61-67	As explained in greater detail below, the data blocks that are to be preserved are first copied into the snapshot storage and a record indicating that the data block has been preserved is updated. Such a record can be stored, for example, in snapshot map <b>52</b> of FIG. 3. New data may then be written to mass storage device <b>20</b> without losing the preserved data blocks.
O4	60C, 61	18/7-11	As will become apparent in the description that follows, snapshot map <b>52</b> of FIG. 3 identifies those data blocks that have changed since a static snapshot was preserved at a particular instant in time.
O5	60D	8/57-62	The amount of data stored in a group or cluster of sectors may also properly be referred to as a data block. If the mass storage means is a RAM or other word or byte addressable storage device, the term data block may be applied to a byte, a word, or multiple word unit of data.
O6	62	13/1-6	During normal operation of a computer system, data is periodically written to or read from attached mass storage means such as mass storage device <b>20</b> . Embodiments within the scope of this invention therefore comprise means for writing data to a mass storage device and means for reading data from a mass storage device.
O7	63,64	3/10-19	A data shadowing program cycles through all the files in a computer network, or through a selected set of critical files and checks the time stamp of each file. If data has been written to the file since the last time the shadowing program checked the file's status, then a copy of the file is sent to a backup system. The backup system receives the data and stores it on tapes or other media. The shadow data is typically more current than data restored from a tape backup, because at least some information is stored during business hours.
O8	75	28/31-38	CRC values are typically used to detect errors in a block of data transmitted across a communication link or stored on a storage device. Cryptographically strong hash functions (also referred to as digests, fingerprints, or message authentication codes) have also been developed to perform a similar function. Any method can be used as long as the digest has a sufficiently high probability of detecting differences between two data blocks.